Microsoft Azure:
Using the Public Cloud to solve the Big Questions

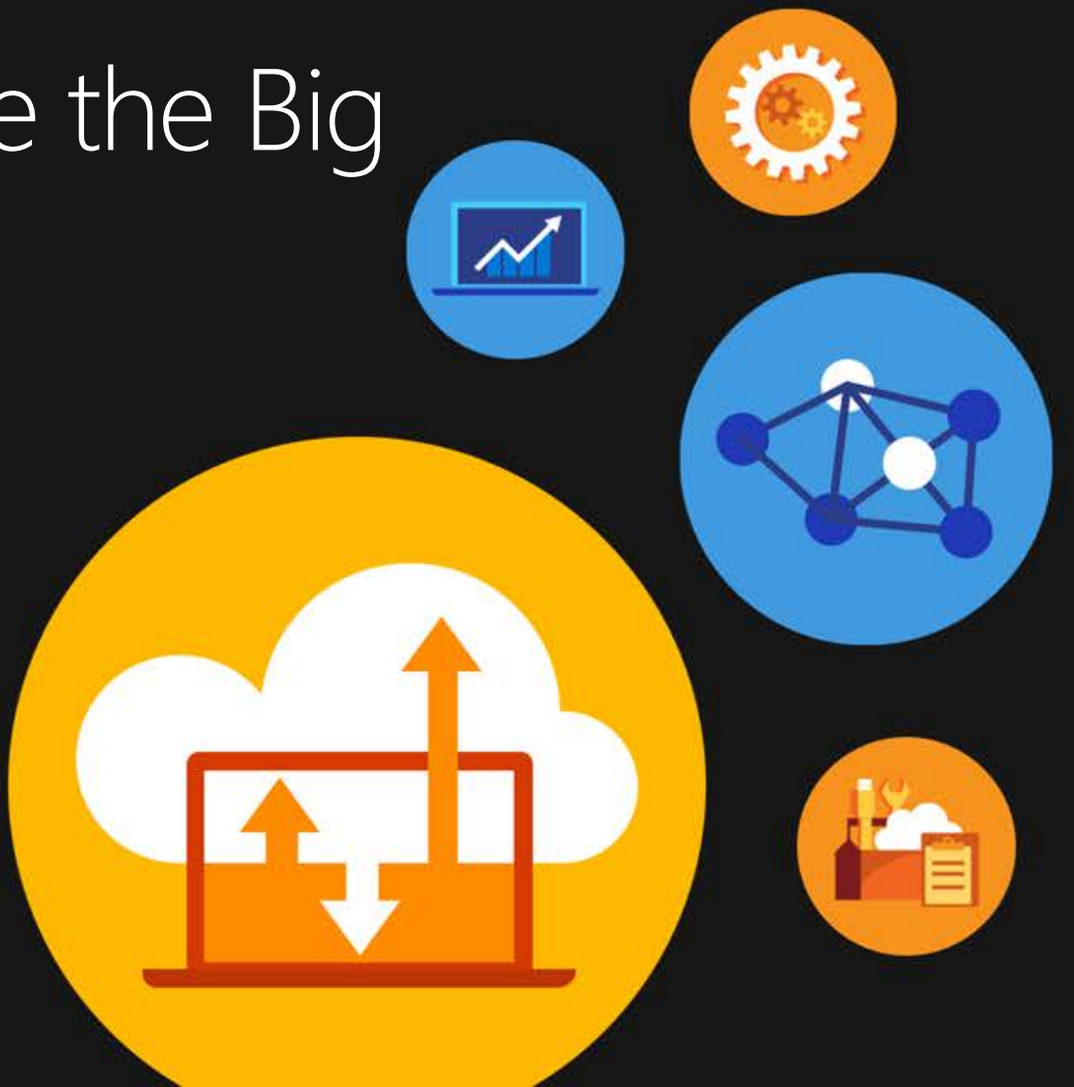

Kent Altena
Global Black Belt TSP, Big Compute
Microsoft
kaltena@microsoft.com
http://microsoft.com/hpc

Microsoft

Microsoft

Introduction to Azure

Hyper-scale infrastructure:

100+ Datacenters across 42 Regions (36 Generally Available) Worldwide.

- Top 3 networks in the world
- 2x AWS, 6x Google DC Regions
- Largest VM in World (large instances, M Series)

Operational

Announced/Not Operational

* Operated by 21Vianet ** Data Stewardship by Deutsche Telekom

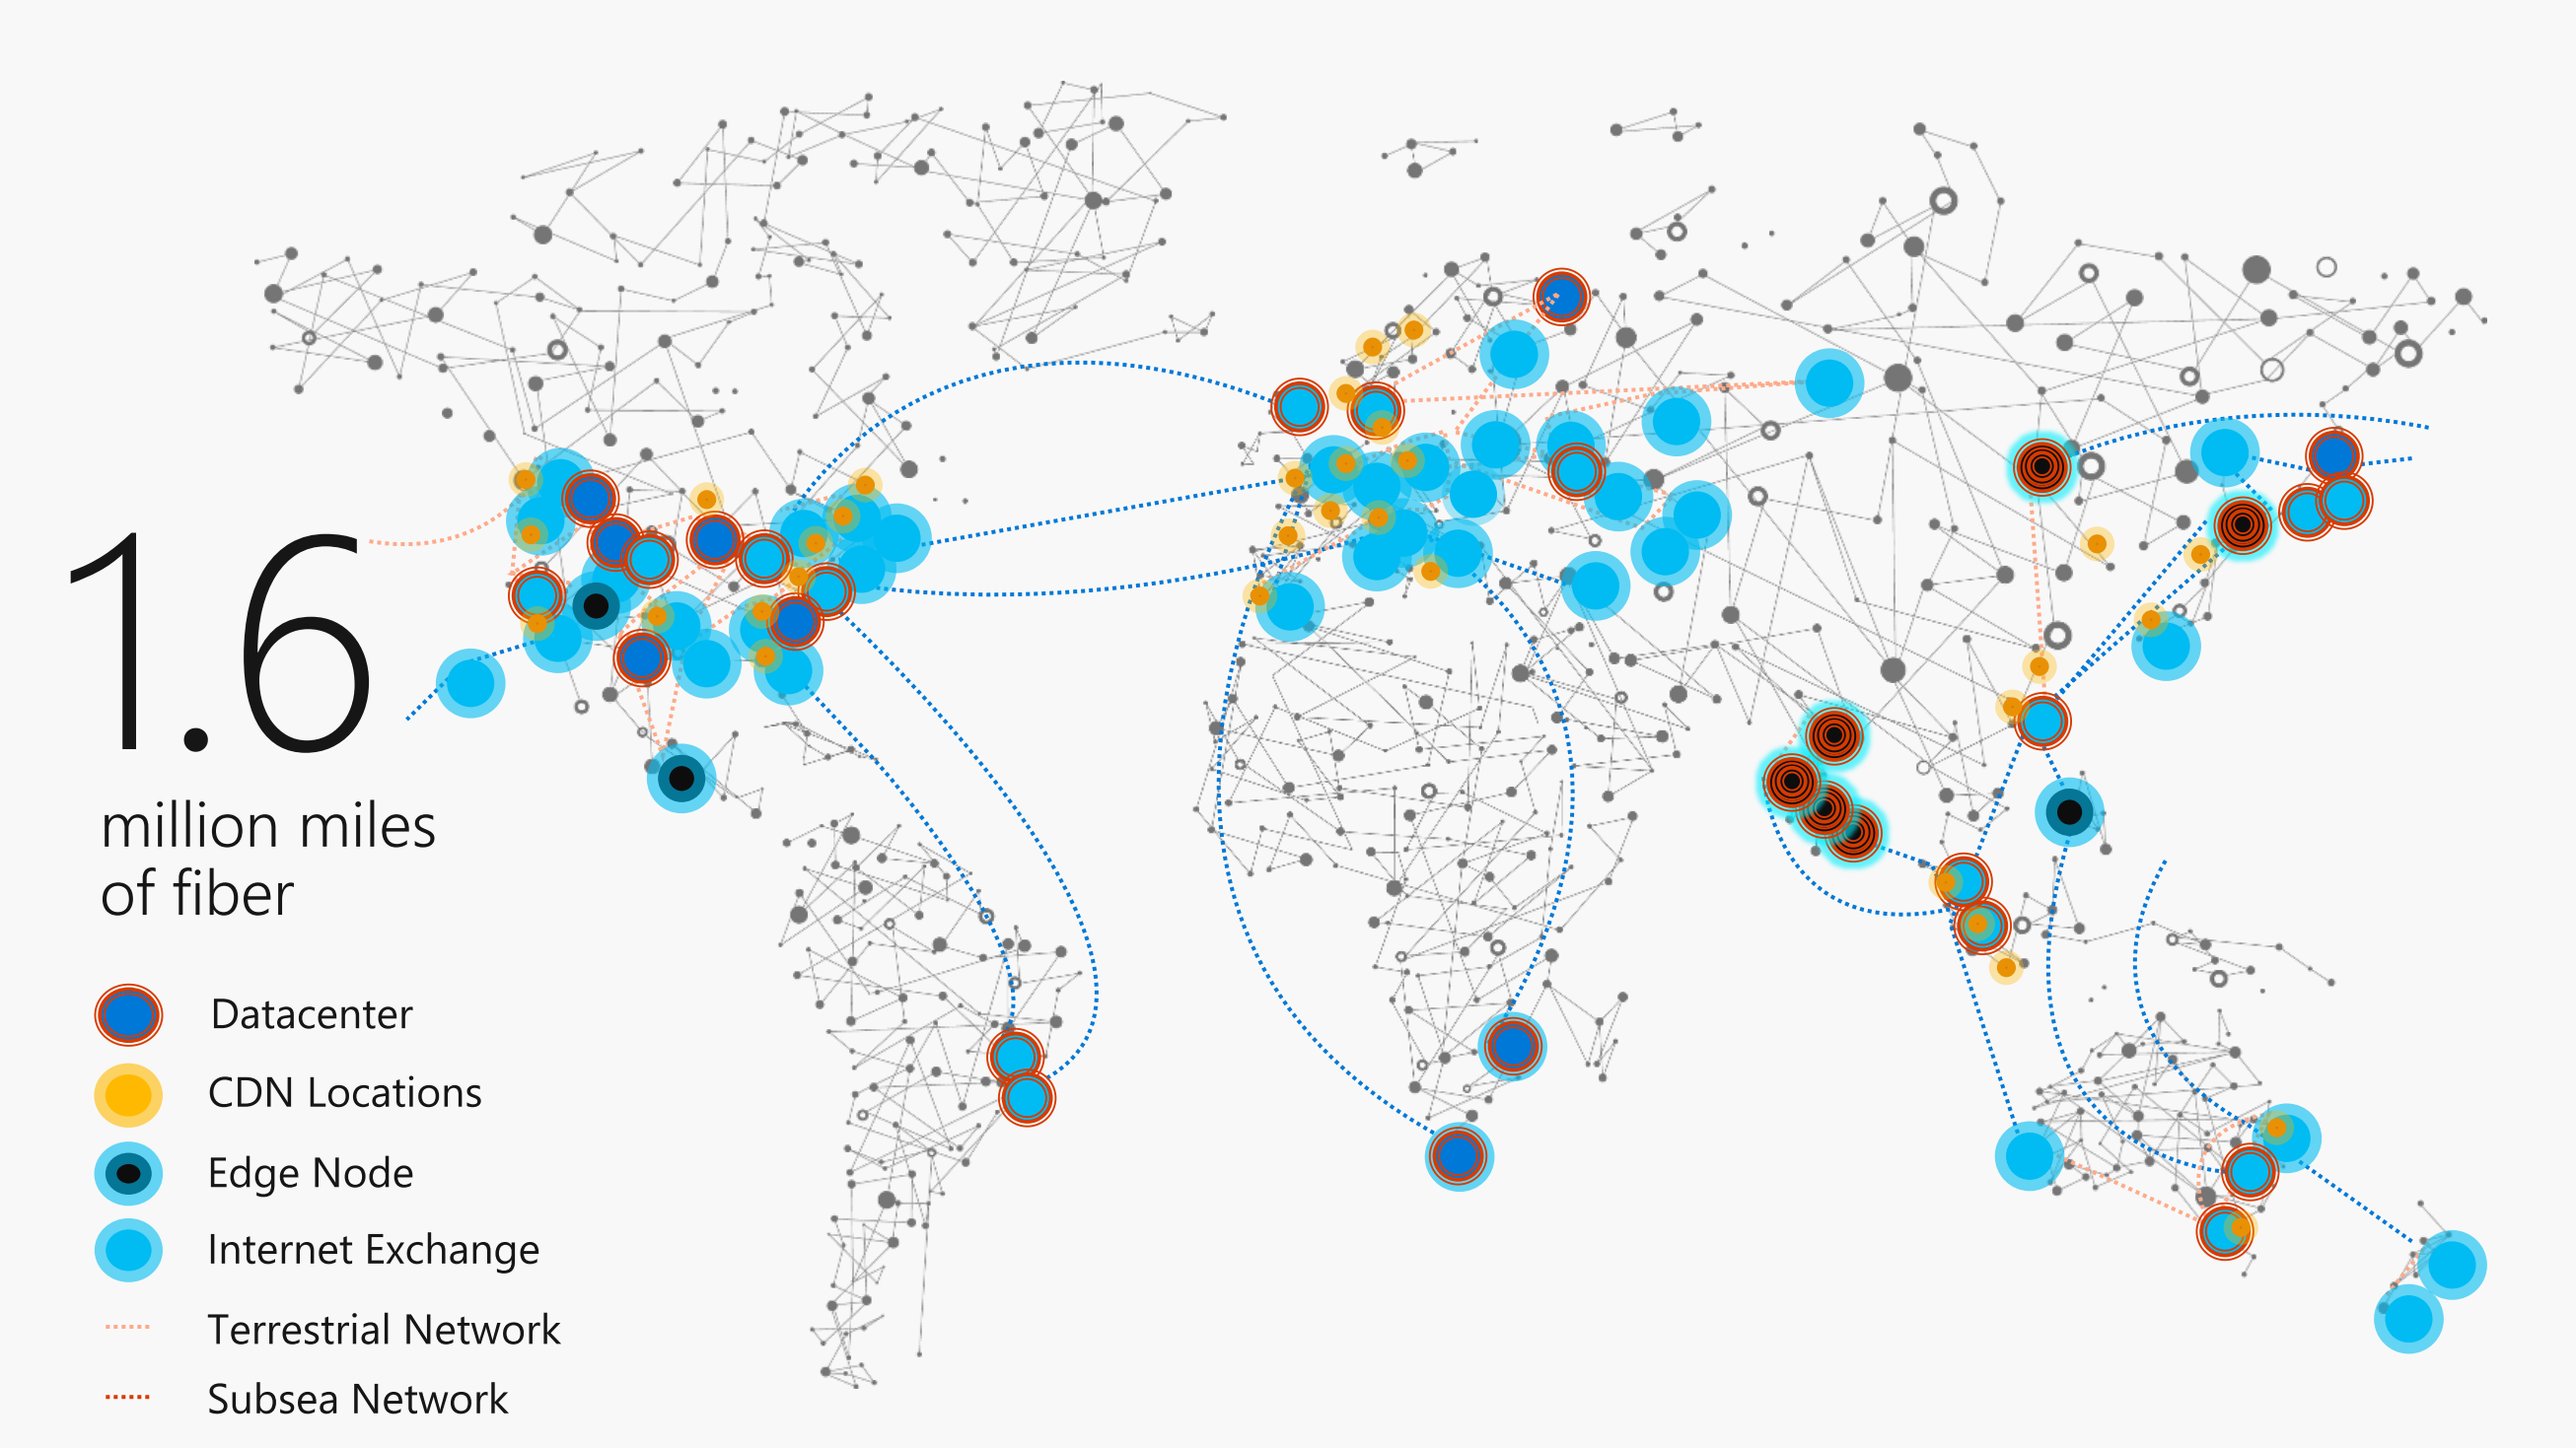
1.6
million miles
of fiber
Datacenter
CDN Locations
Edge Node
Internet Exchange
Terrestrial Network
Subsea Network

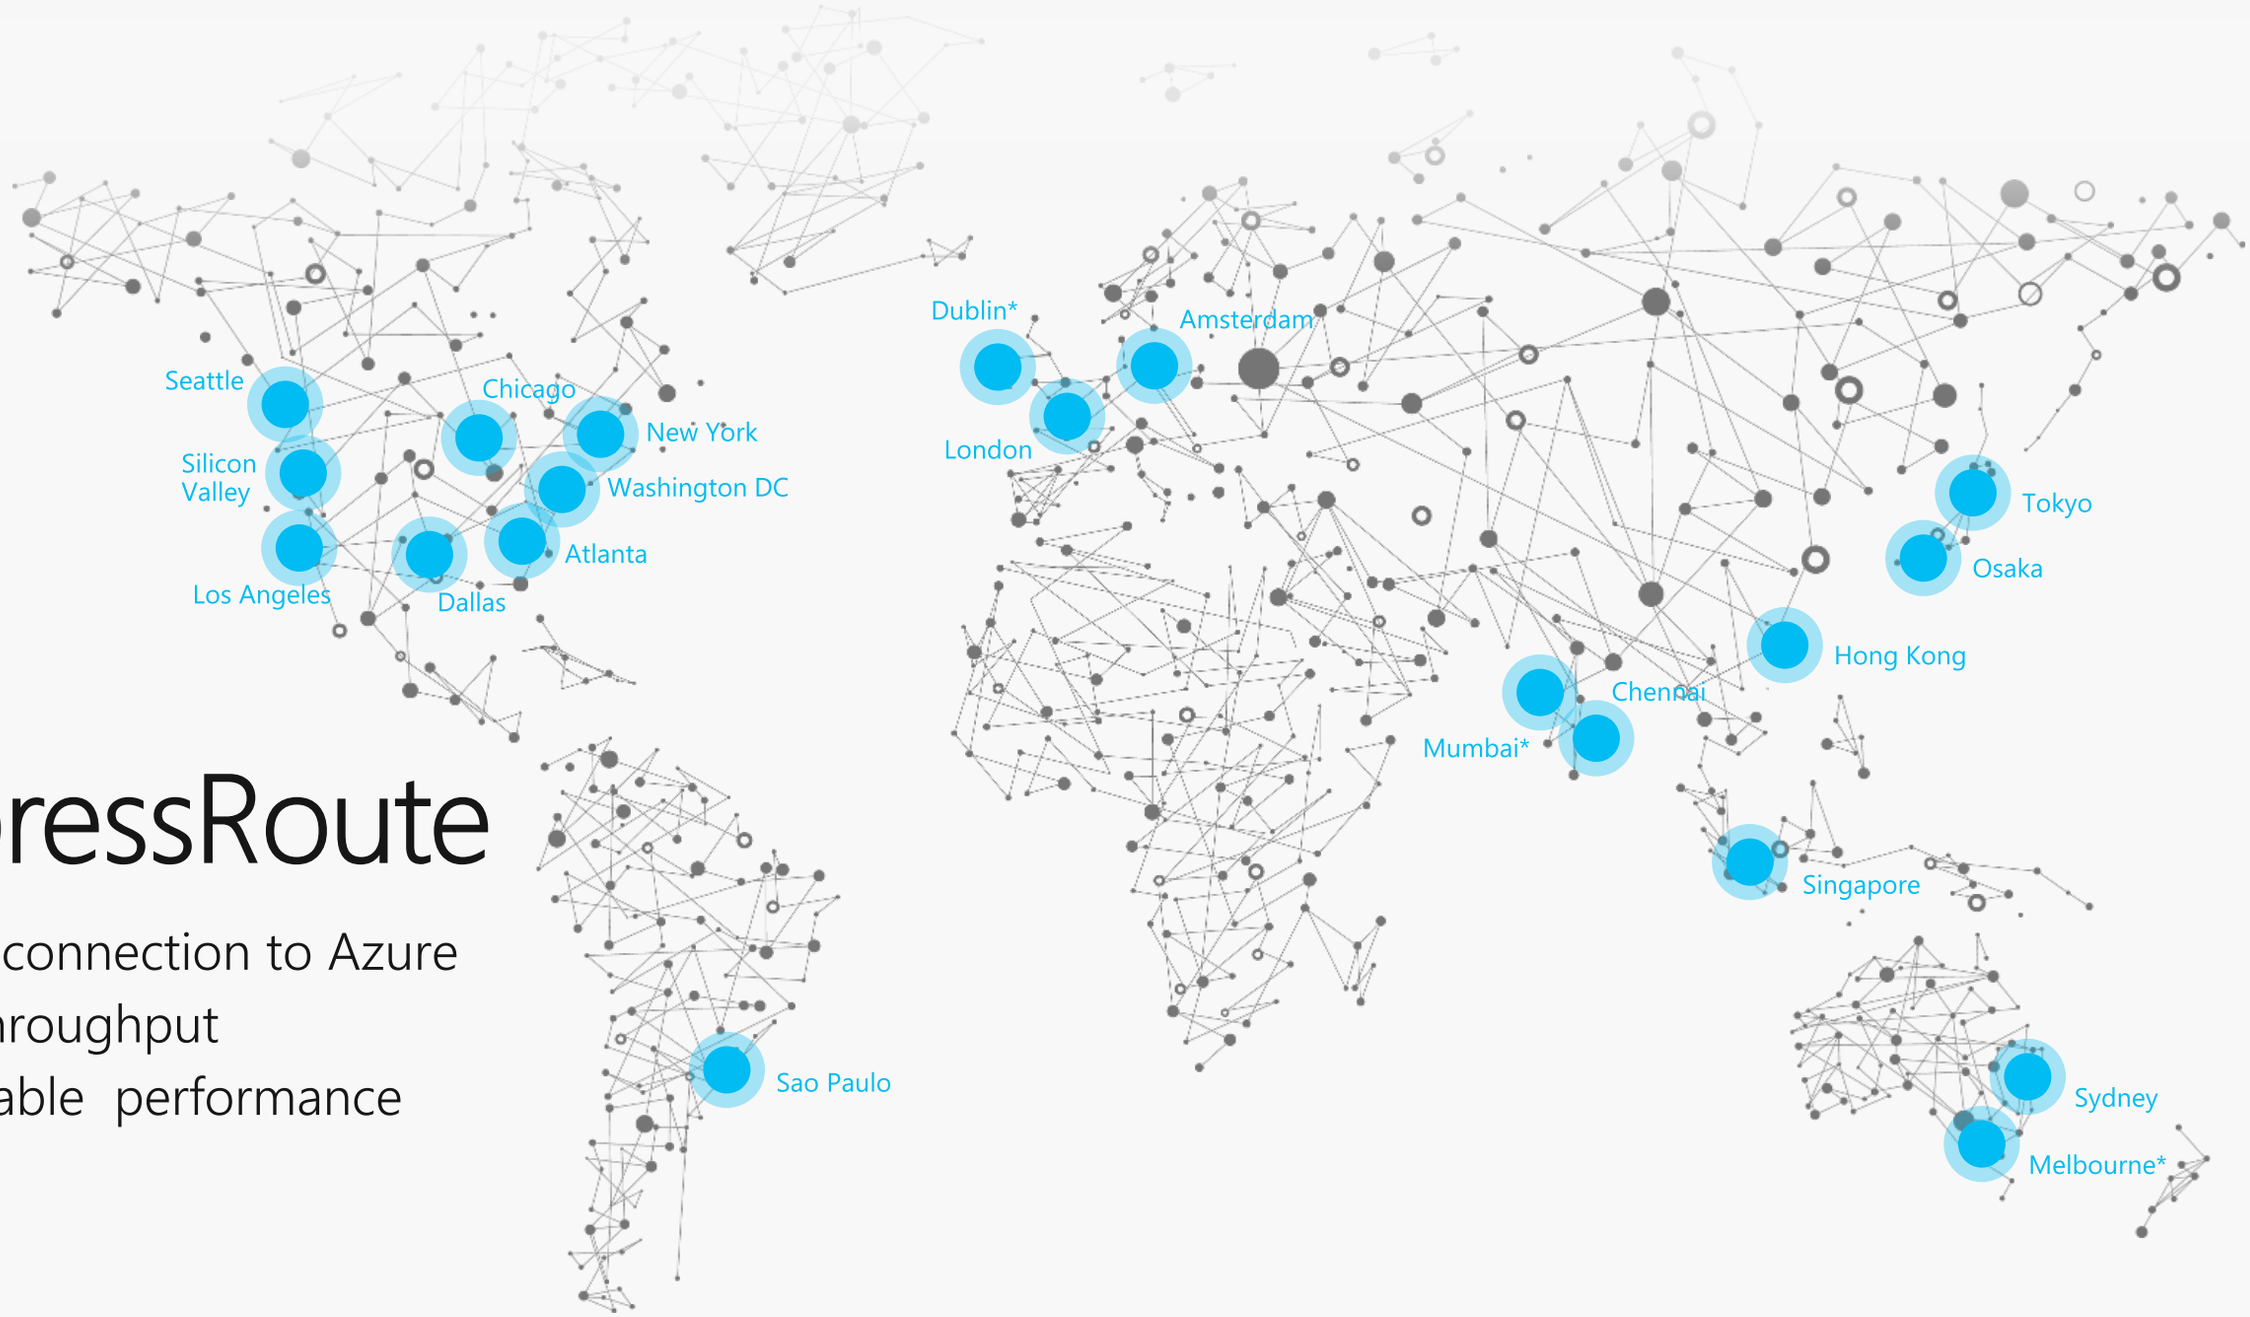

ExpressRoute

Private connection to Azure

High-throughput

Predictable performance

Comply with industry and regional requirements

GLOBAL

- ISO 27001
- ISO 27018
- ISO 27017
- ISO 22301
- SOC 1 Type 2
- SOC 2 Type 2
- SOC 3
- CSA STAR Self-Assessment
- CSA STAR Certification
- CSA STAR Attestation

US GOV

- Moderate JAB P-ATO
- High JAB P-ATO
- DoD DISA SRG Level 2
- DoD DISA SRG Level 4
- DoD DISA SRG Level 5
- SP 800-171
- FIPS 140-2
- Section 508 VPAT
- ITAR
- CJIS
- IRS 1075

INDUSTRY

- PCI DSS Level 1
- CDSA
- MPAA
- FACT UK
- Shared Assessments
- FISC Japan
- HIPAA / HITECH Act
- HITRUST
- GxP 21 CFR Part 11
- MARS-E
- IG Toolkit UK
- FERPA
- GLBA
- FFIEC

REGIONAL

- Argentina PDPA
- EU Model Clauses
- UK G-Cloud
- China DJCP
- China GB 18030
- China TRUCS
- Singapore MTCS
- Australia IRAP/CCSL
- New Zealand GCIO
- Japan My Number Act
- ENISA IAF
- Japan CS Mark Gold
- Spain ENS
- Spain DPA
- India MeitY
- Canada Privacy Laws
- Privacy Shield
- Germany IT Grundschutz workbook

A quick overview of HPC

What is HPC all about?

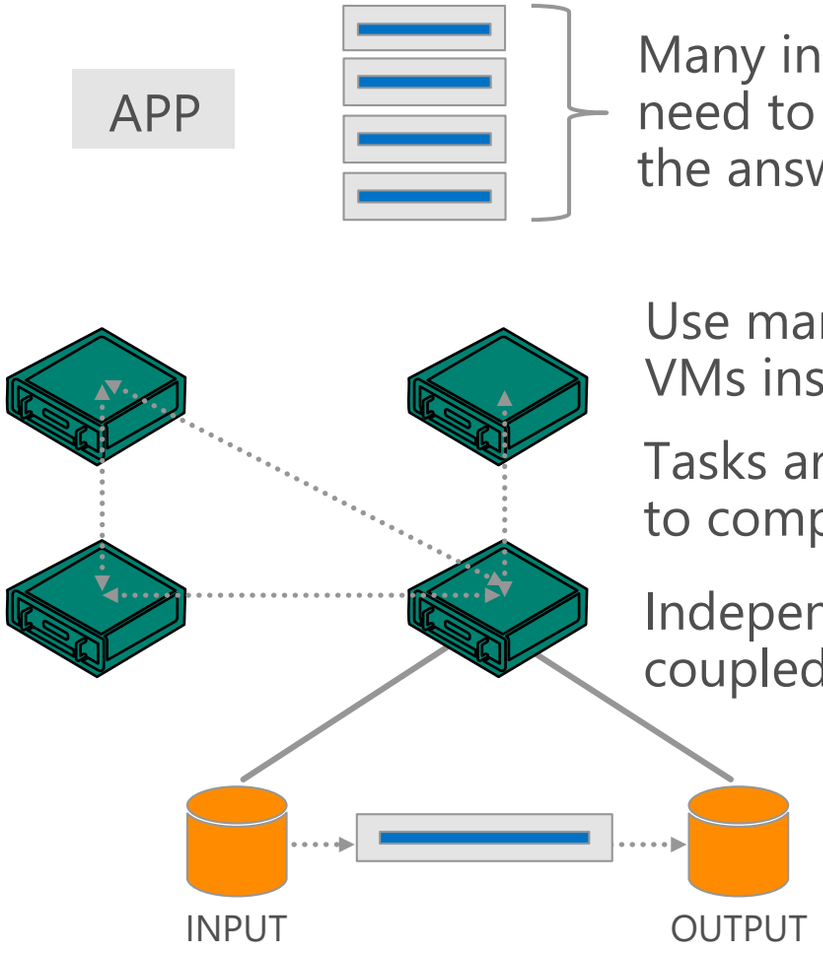

Data is read, computed and written/stored

Uses:

Science and research

Genomics & bioinformatics

Climate modeling

Hydrological analysis

Image analysis & processing

Video & audio transcoding

Engineering stress analysis

Automobile crash simulation

Financial risk analysis

Oil and gas production

… and much more!

Two main types of HPC workloads

- Embarrassingly parallel:
- Nodes don't need to talk to each other, or very little cross-node communication
- Usually a parameter sweep, a job splitting, or a search/comparison through data
- Examples: Monte Carlo simulations, image/video rendering, genetic algorithms, sequence matching
- →Great workload for the cloud!

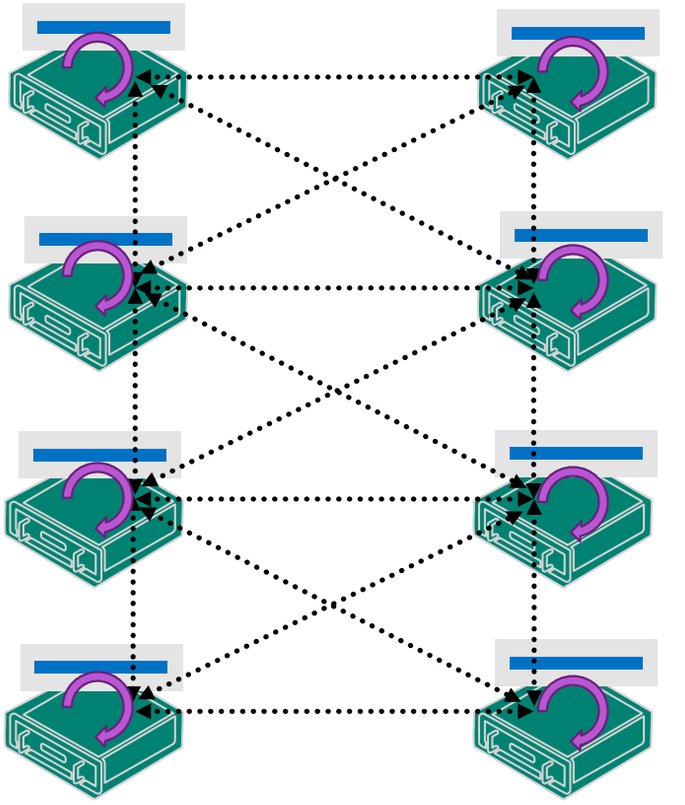

Tightly coupled:

- Nodes need to talk to each other constantly
- Requires a fast interconnection network (low latency and high throughput)
- Examples: automotive crash simulation, fluid dynamics, climate modeling, reservoir simulation, manufacturing modeling

→More challenging, but already possible on Azure!

Mission

Our mission is to enable researchers, engineers, analysts, designers, developers, and data scientists to achieve radically better results and faster answers to complex problems, by making it easy to do simulation and parallel computing in the cloud, at hyper-scale.

Azure empowers these technical experts to think at 10x or 100x the current scale, to work faster, better, and in new ways that before the cloud we could only dream about.

Azure = True HPC on the cloud

- Performance-tuned CPU and BIOS provide supercomputing-class VM performance to run jobs faster and reduce costs – no hidden hyper-threading!
- Engineering in the platform enables Azure to deliver speeds comparable to (and sometimes better than) bare-metal hardware on premises
- Azure is the only large cloud provider with InfiniBand RDMA networking and NVIDIA GRID license included for GPU workloads
- Support for both Linux and Windows workloads
- Multiple deployment options allow you to leverage existing technology investments

Really? True HPC on the cloud?

From a 3rd party, independent research published on 9, February, 2017 :

Comparative benchmarking of cloud computing vendors with High Performance Linpack

Mohammad Mohammadi, Timur Bazhirov, Exabyte Inc., San Francisco, California 94103, USA

"We found Microsoft Azure to deliver the best results, and demonstrated that the performance per single computing core on public cloud to be comparable to modern traditional supercomputing systems. Based on our findings we suggest that the concept of high performance computing in the cloud is ready for a widespread adoption…"

The media's reaction was quick:

"The actual TOP500 run across the entire machine achieved a Linpack yield of more than 64 percent of peak, which is fairly typically of an HPC cluster with a high-performance network. The Azure H16 in this test had a 67 percent Linpack yield."
-Top500.org

performed best, with
node [...] The Edison
e to 32 nodes."
-The Register

ranked su
-Fudzilla

-The Stack

What would you do with 100x the scale?

- Do more with hyper-scale:
- Service more users
- Run more projects
- Get results faster

- Run larger simulations
- Explore new insights (e.g., “What if?”)

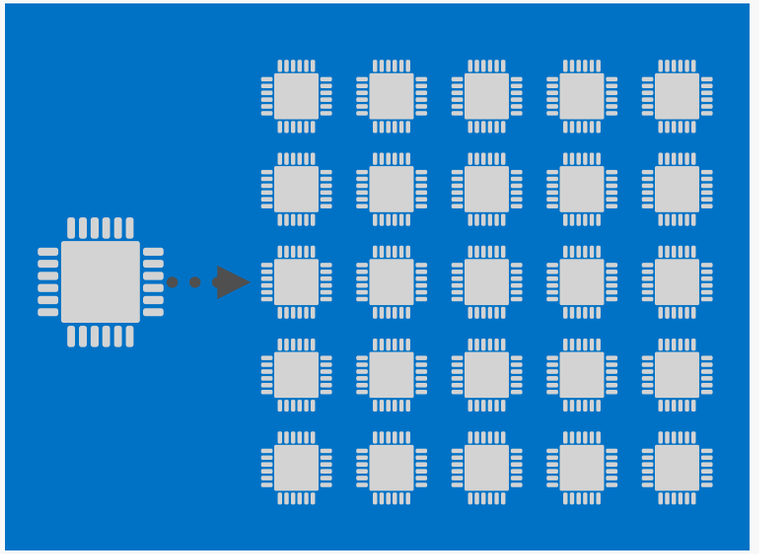

Remove current limitations:

- Modify more parameters
- Analyze more complex models
- Visualize larger results

- Run more iterations
- Generate higher fidelity results
- Simulate longer periods of time

Azure HPC building blocks

- Hardware capabilities: a range of infrastructure and platform offerings, including HPC-class infrastructure, to match the requirements of any type of parallel workload.

- Reference Architectures / Services: a combination of advanced tools, applications, and cloud services to help you run HPC workloads on Azure, at any scale.

Hardware Capabilities for HPC

Hardware capabilities

VM sizes

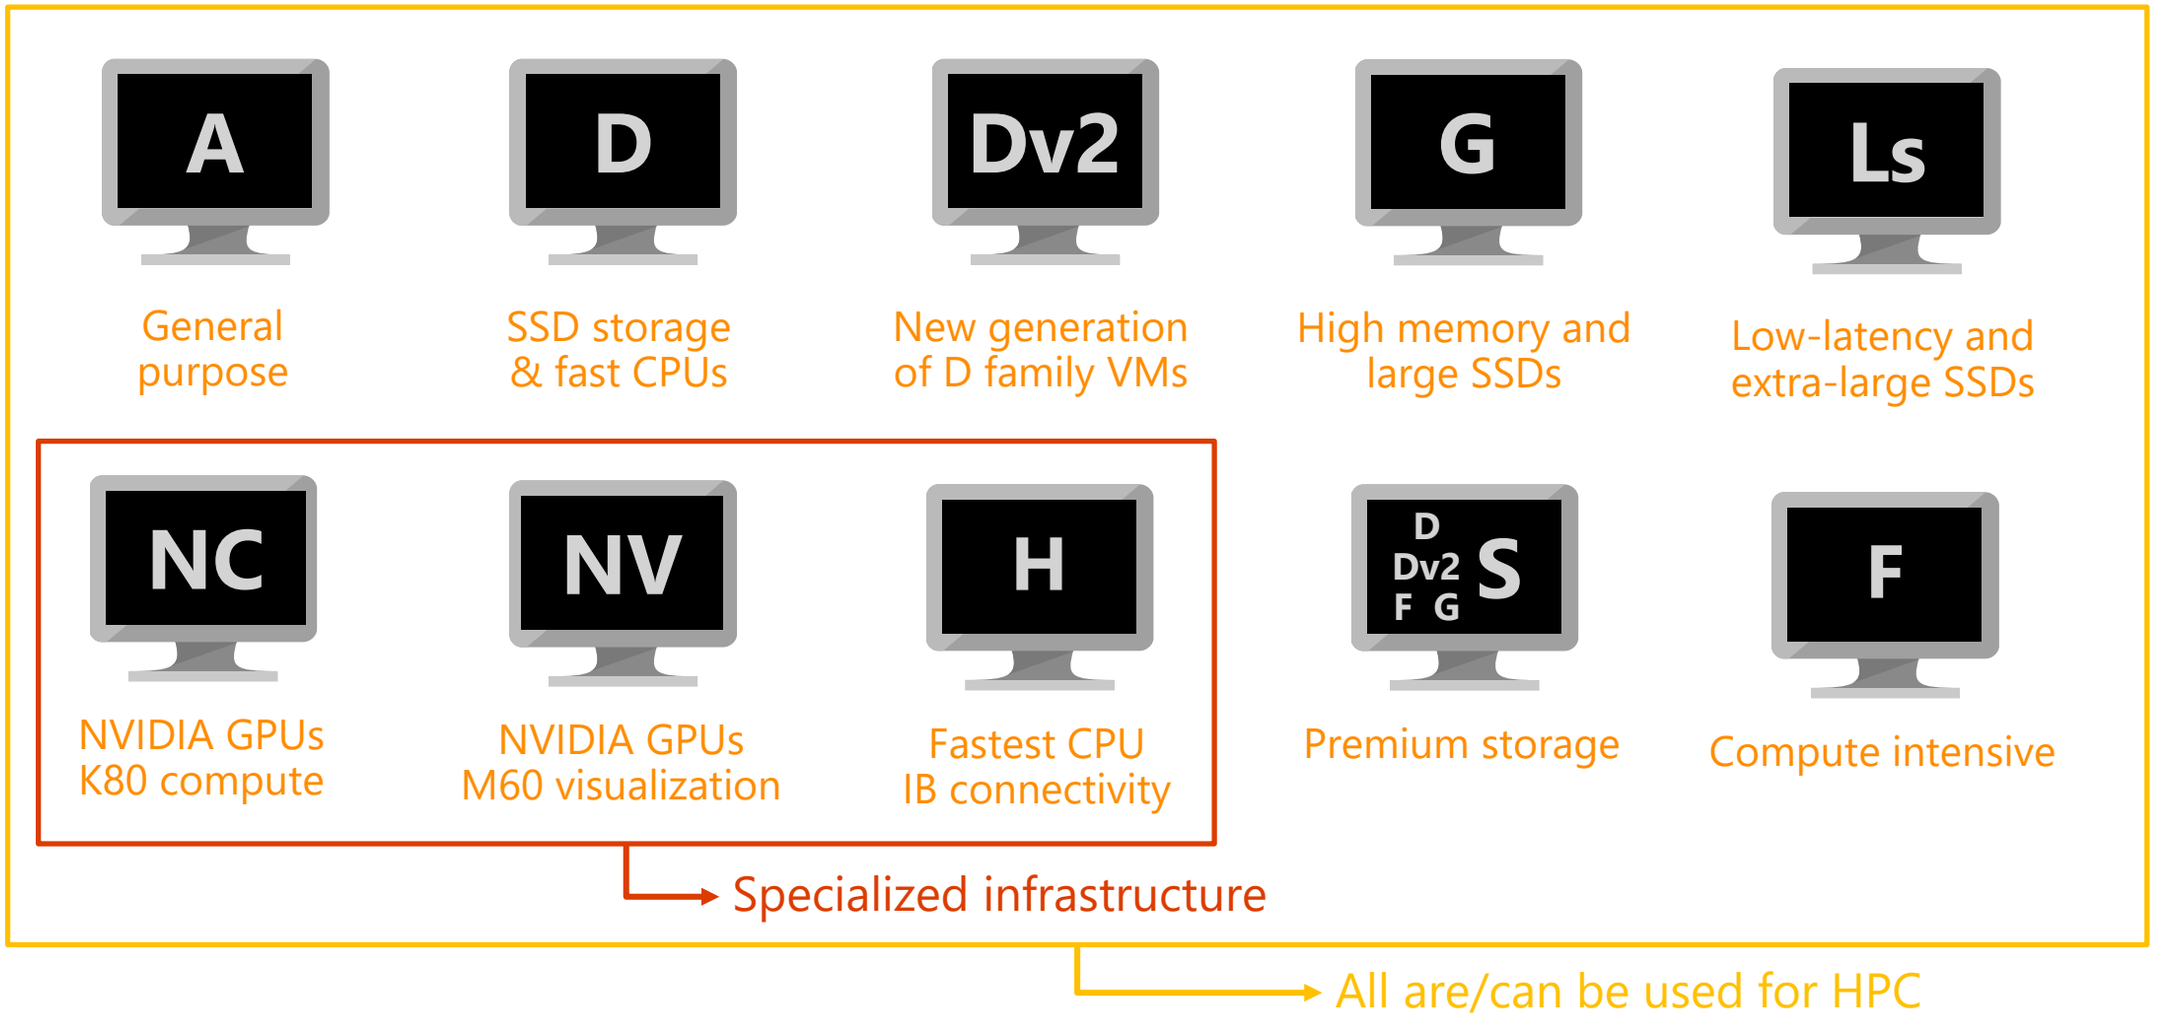

Hardware capabilities

Specialized infrastructure

- A8-A9 and A10-A11 are performance-tuned VMs that provide bare-metal results
- 8 or 16 CPU cores per VM
- A8-A9 with QDR InfiniBand + RDMA on Windows and Linux
  - < 3 µsecs latency
  - 32 Gbps throughput (actual)
- Fully non-blocking topology to ensure max bandwidth
- E5-2670 Sandy Bridge processor

- H series
- Fastest VMs in the public cloud
- 8 or 16 CPU cores per VM
- H16r and H16mr with FDR InfiniBand + RDMA on Windows and Linux
  - < 3 µsecs latency
  - 54 Gbps throughput (actual)
- Fully non-blocking topology
- E5-2667 v3 Haswell processor
- Up to 224 GB DDR 4 memory
- Up to 2 TB of local SSD

- NC and NV series
- Up to 4 GPUs per VM
- Up to 24 CPU cores per VM
- NC24r with FDR InfiniBand + RDMA on Windows and Linux
  - < 3 µsecs latency
  - 54 Gbps throughput (actual)
- Fully non-blocking topology
- E5-2667 v3 Haswell processor
- Up to 224 GB DDR 4 memory
- Up to 2 TB of local SSD

What's the difference between CPU and GPU?

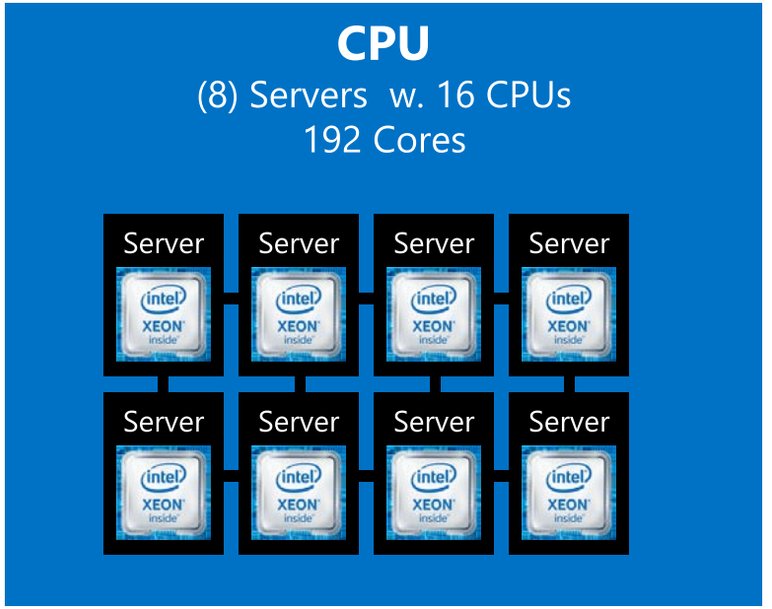

A CPU is composed of just few cores with lots of cache memory that can handle a few software threads at a time.

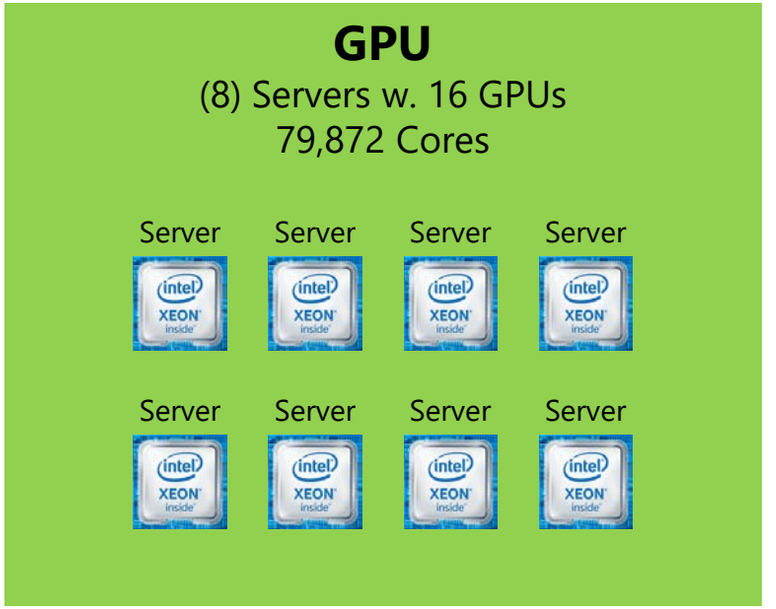

A GPU is composed of hundreds of cores that can handle thousands of threads simultaneously – accelerating some software by 100x vs a CPU.

Optimized to take huge batches of data and performing the same operation repeatedly very quickly. More power- and cost-efficient than a CPU.

New: Low-Priority VM's

Significantly lower priced compute

- Up to 80% discount compared to on-demand price - fixed price
- All Batch VM sizes and regions
- Uses surplus capacity; availability could vary; VMs could be preempted

Suitable workloads

- Distributed parallel jobs - many discrete tasks, interrupt tolerant, shorter task execution times, flexible job completion time
- e.g. Dev, test – regression, scale, load

Value

- Get work done for lower cost, faster, or do more for same price

Reference Architectures for HPC in the Cloud

Software & services: Cloud burst

Add cloud resources to your cluster, on demand

- Burst to cloud to add resources when needed
- Continue using your existing on-premises infrastructure
- Run workloads on Windows and Linux, on Azure and on-premises
- Cover peaks in demand or special projects
- Pay only for what you use
- Microsoft HPC Pack, Univa Grid Engine, and Altair PBS Pro already have this capability (and more are coming!)

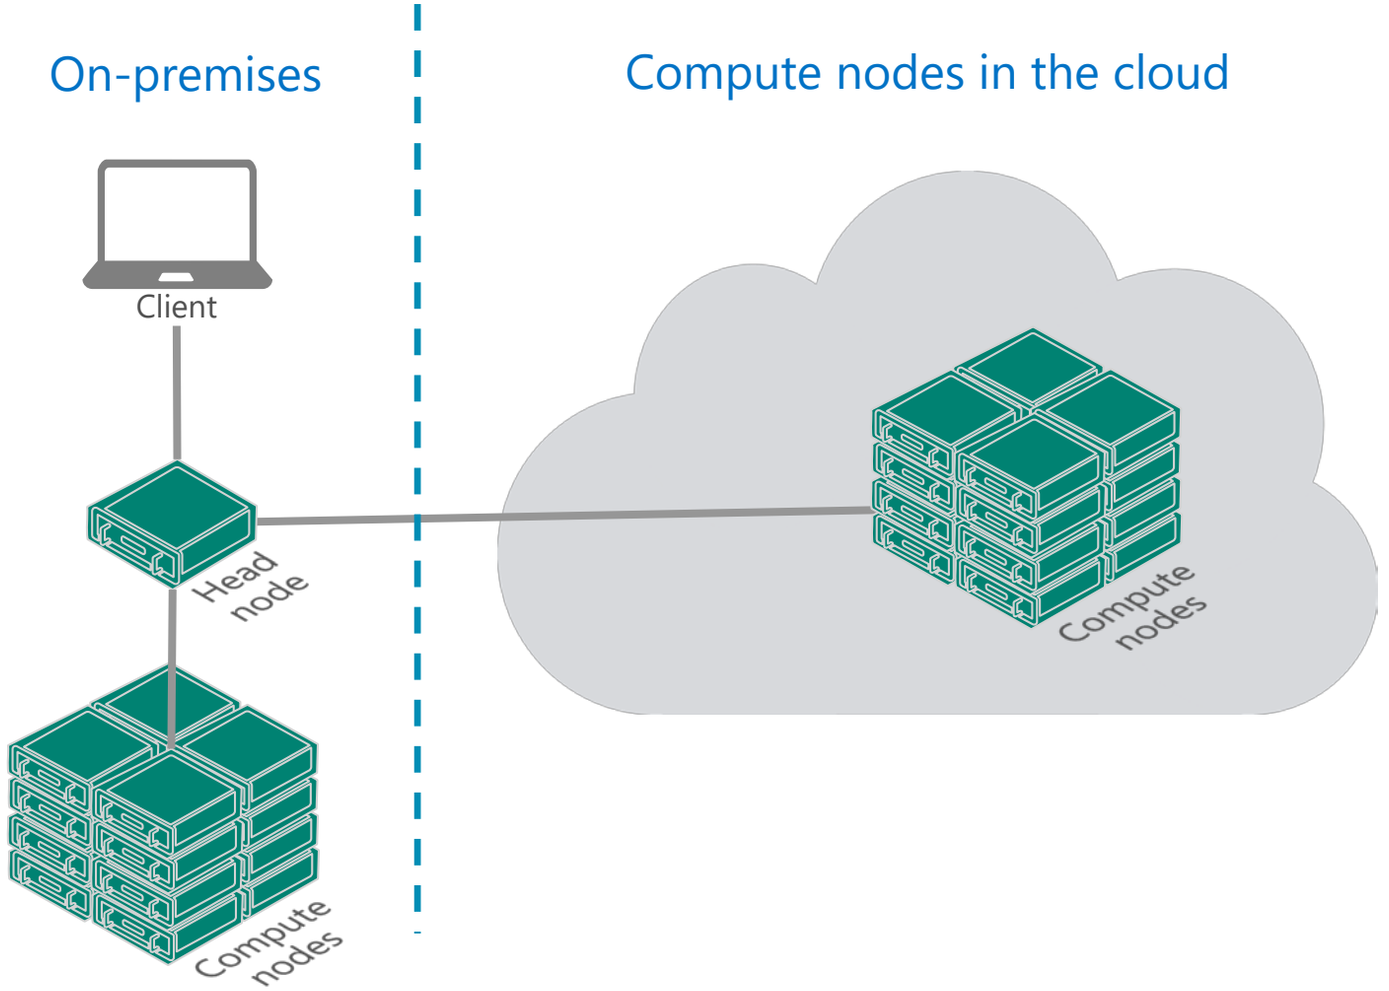

Software & services: Cluster on the cloud

Provision one (or more) new clusters in minutes

- Deploy a complete HPC cluster, all in the cloud, in just minutes
- Use templates, scripts, and images to quickly deploy at any scale
- Use your current HPC scheduler
- Shift existing applications, scripts and tools to cloud
- Deploy as many clusters as you need!

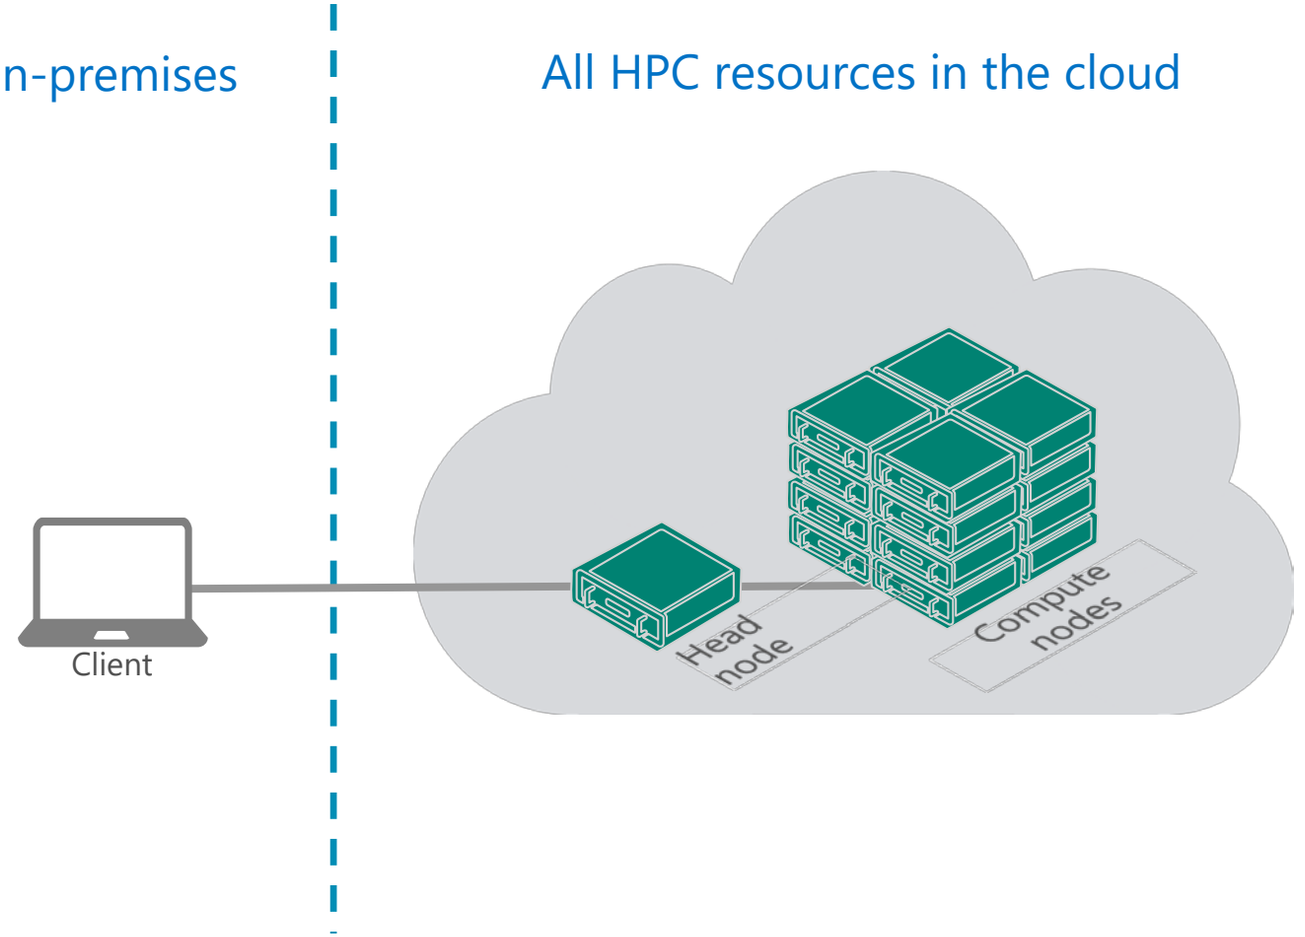

Software & services: HPC as a service

Run at scale directly from your application

- Integrate with Azure Batch, directly from a client application (GUI or CLI) or online portal
- Batch abstracts resource management and scheduling completely
- Supports small to extremely large deployments and can deploy any VM size
- Provides auto-scaling and stopping of resources
- Run HPC jobs at scale on Docker containers
- Using Batch is free, you only pay for the underlying compute!

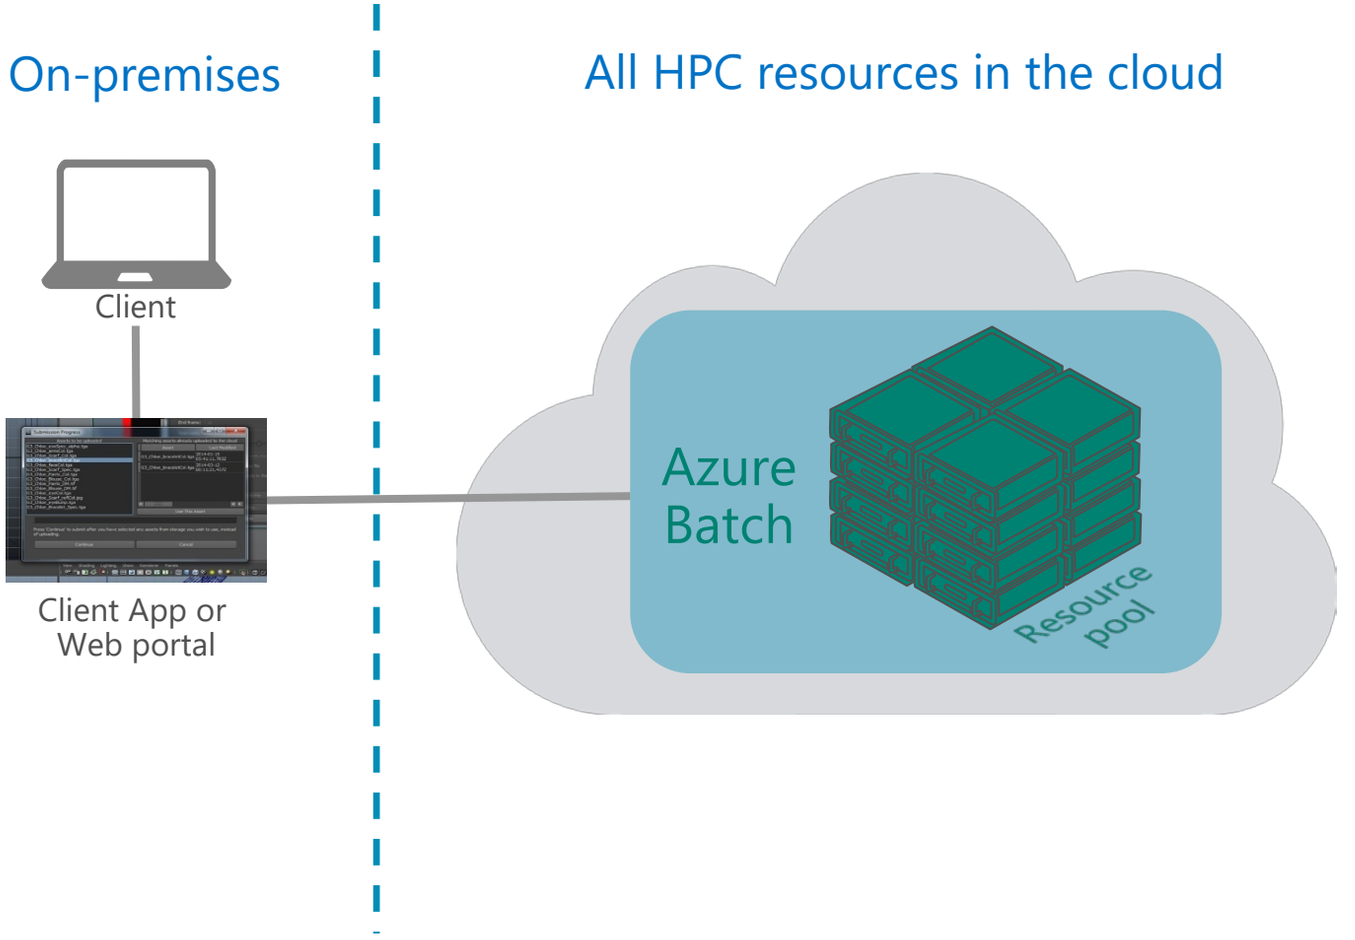

Why Batch? Because customers wanted something to make HPC easier on the cloud

"I just want to click Run and get my results"

"I don't want to pay for compute nodes to run 24/7, I only need them to run one job"

"We didn't get approved for a refresh of our 2000 core cluster and are now looking at cloud options…but we first need to learn how to manage HPC in the cloud"

"We are a research team and do not know how to manage infrastructure"

"Managing infrastructure is painful and time-consuming"

"I want to run my models from my workstation and not have to move my data to a cluster and run the models from there"

Building an HPC solution on (any) cloud

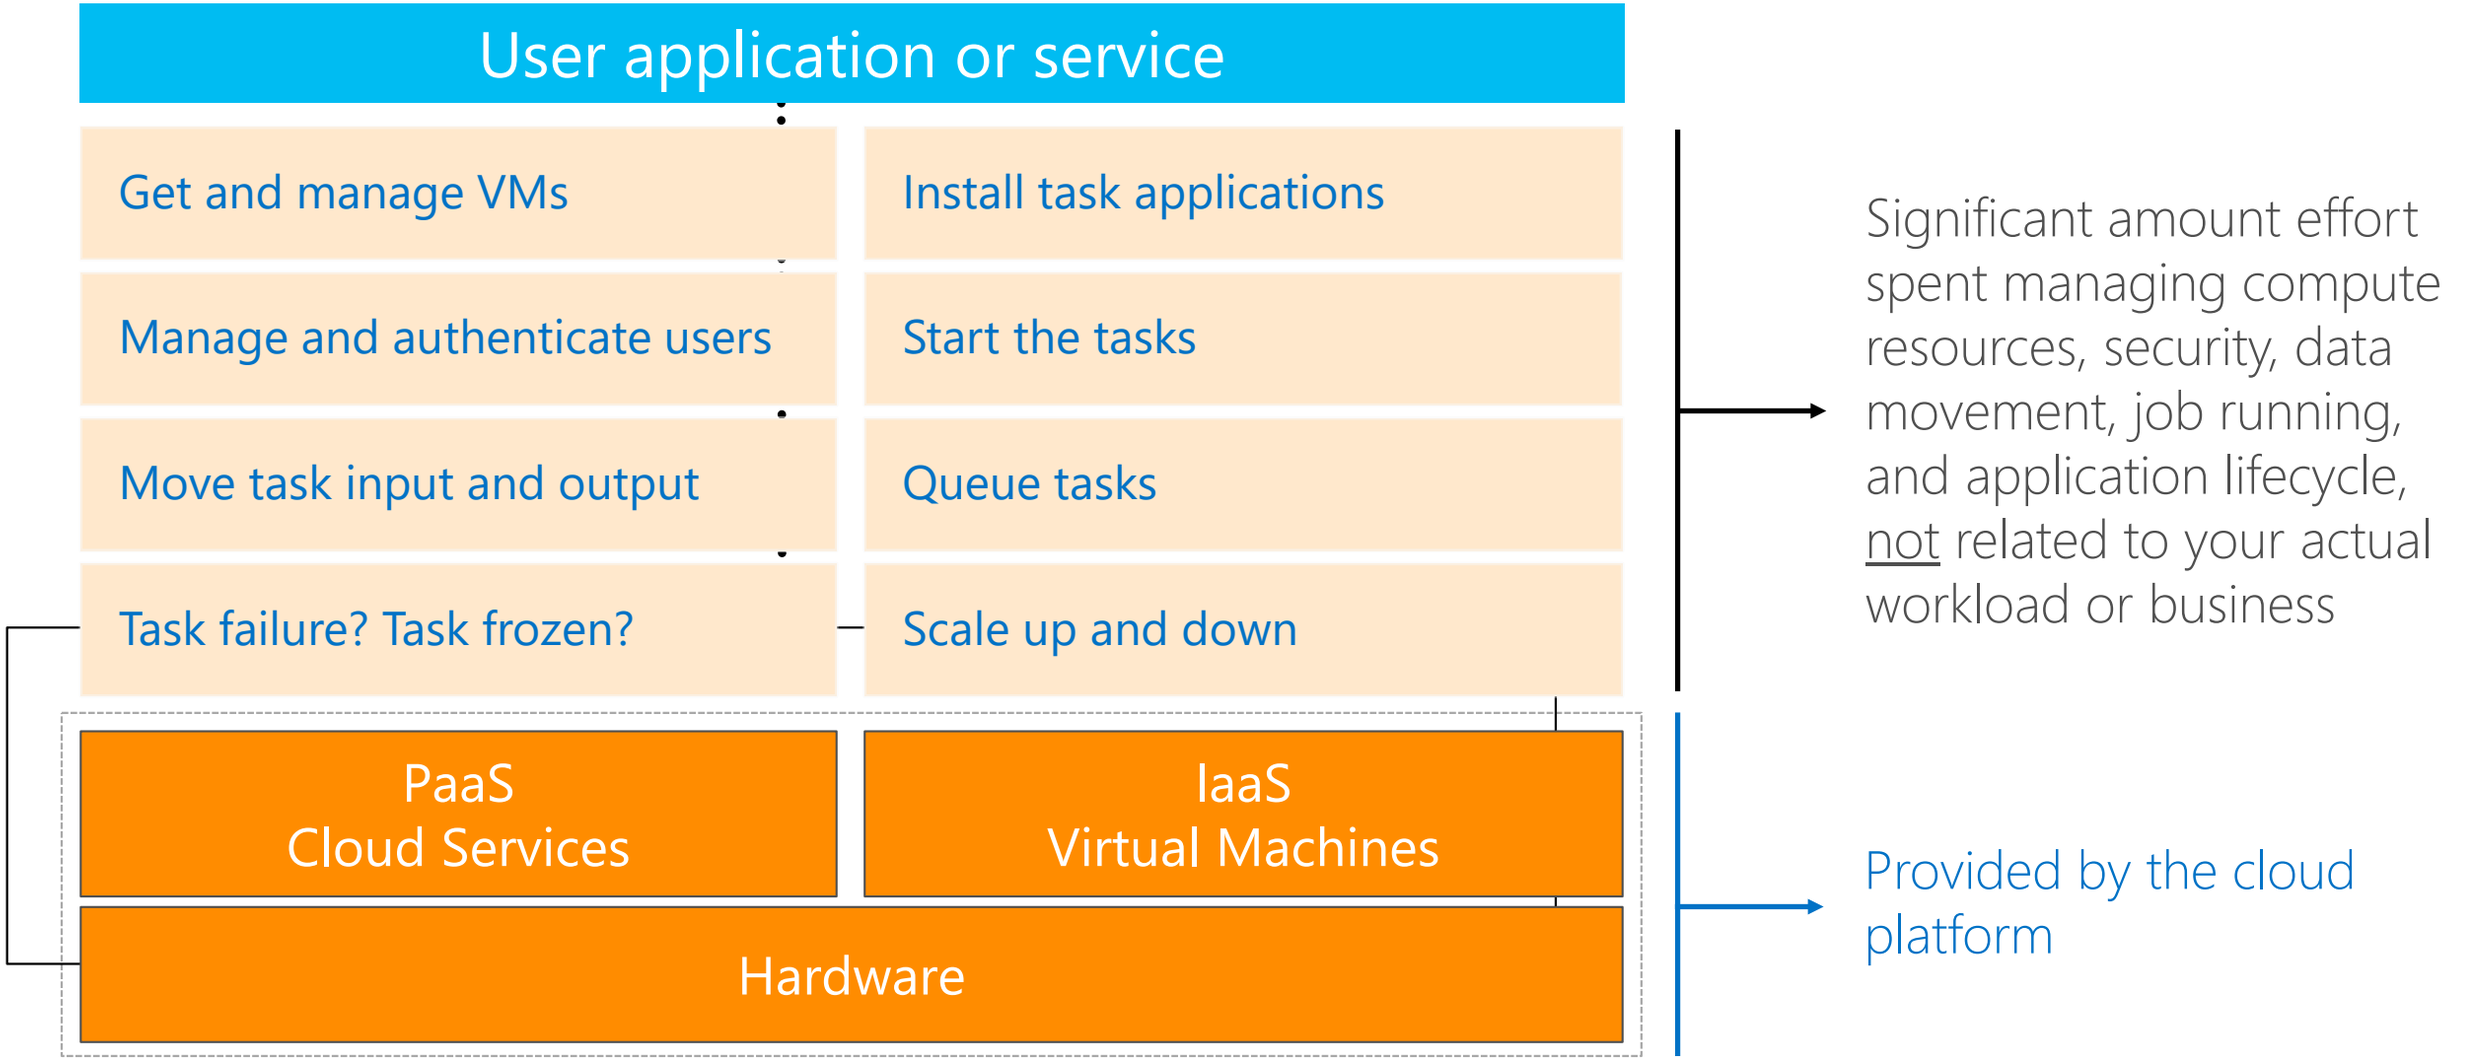

Azure Batch: HPC as a Service

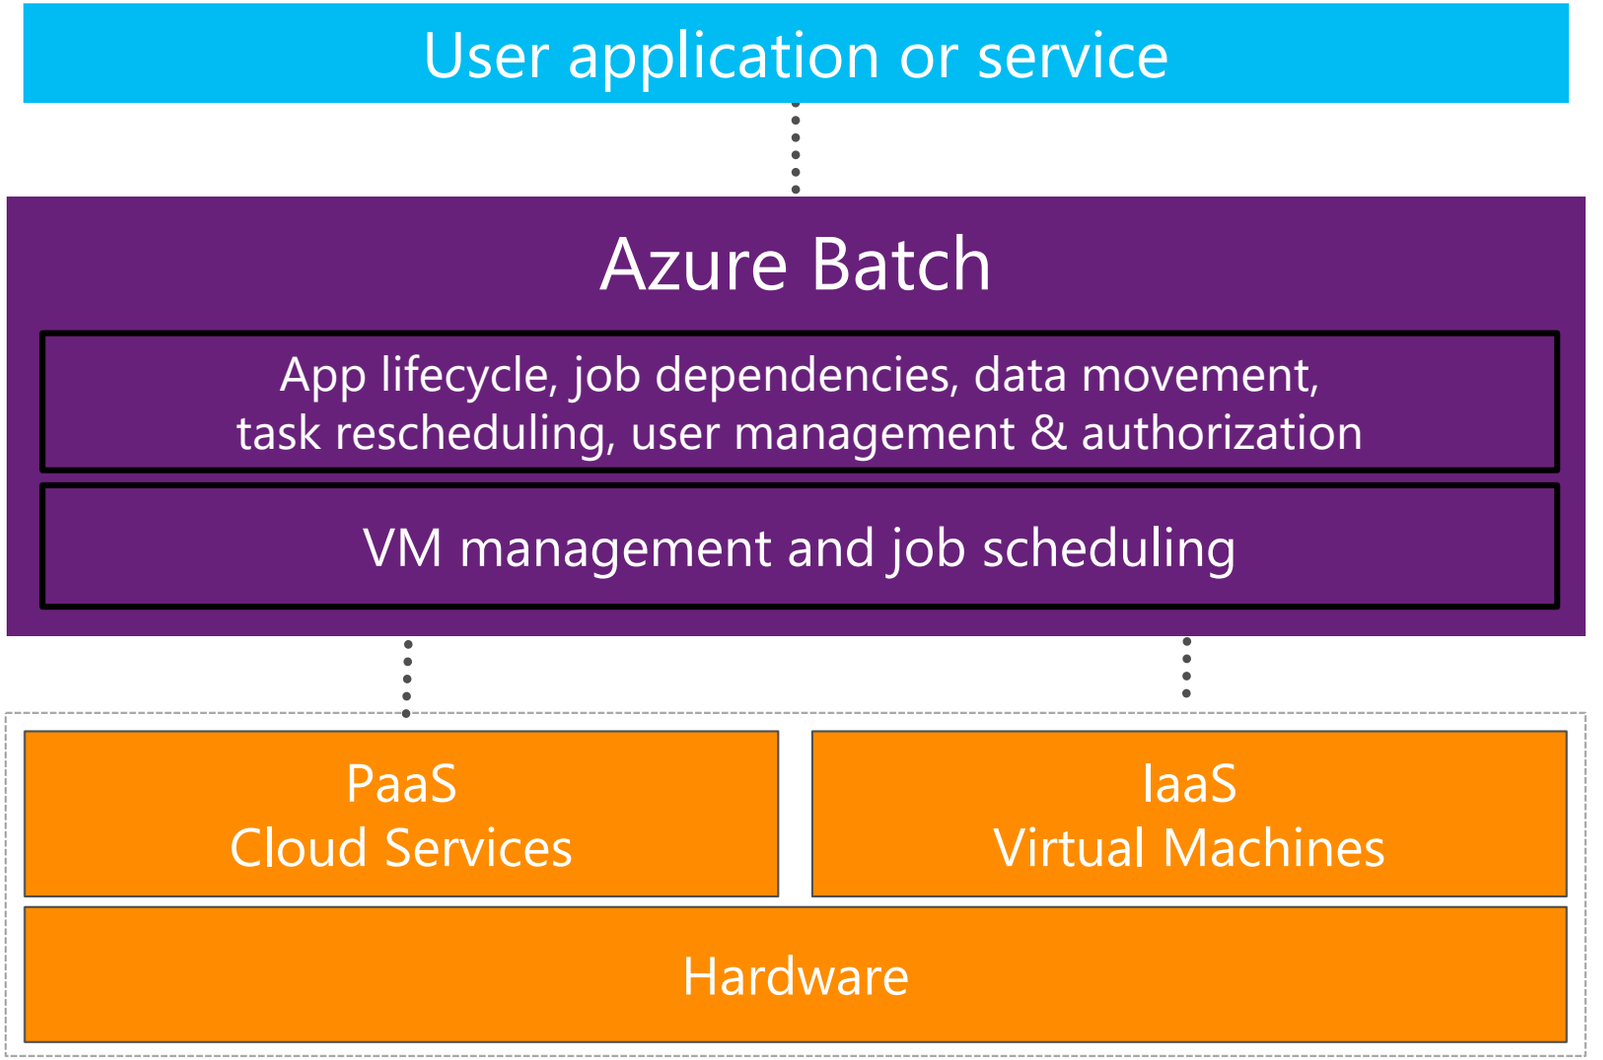

Provided by the cloud platform

- Don’t worry about the “plumbing”
- Focus on the workload/app
- Access higher-level capabilities
- Minimize the required cloud or Azure experience

Cloud-enable applications

Developer

Client plugins

Application & files

Job splitter & task processor

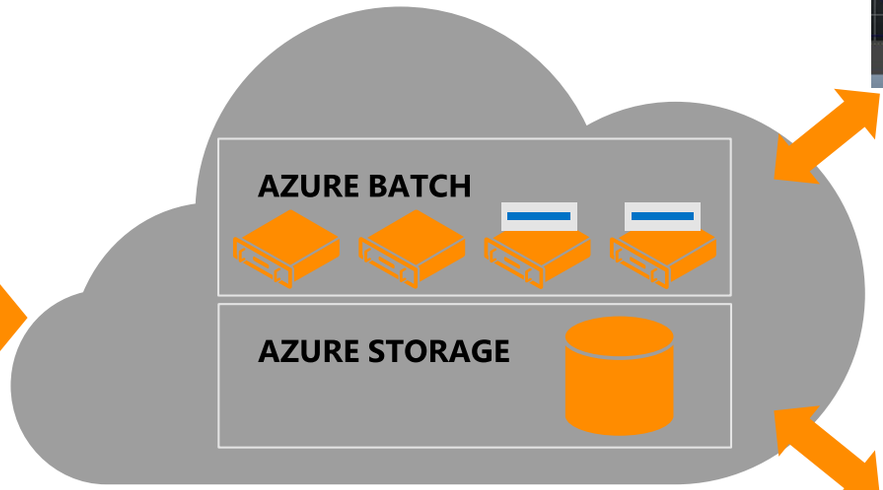

Researcher

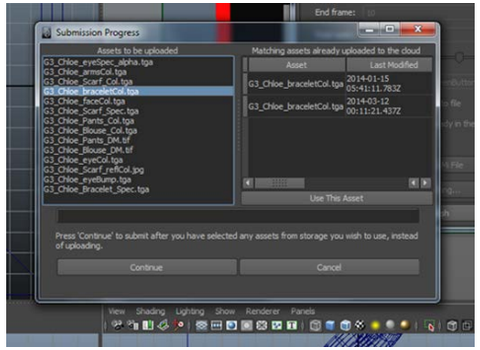

- Submit jobs from client
- Sync input/output data
- Monitor via plugin and/or portal

Ops/Admin

- User & group management
- Role-based access
- Job & VM monitoring
- Usage reporting
- Quota management
- Troubleshooting

Demo:
Using R-Studio to use Azure Batch

Case studies

National Oceanic and Atmospheric Administration

On-Premises: Low Priority / Too Slow

NOAA Researchers have to wait in line for the on-premises Super Computer.

Researchers studying the effect of windmills on migratory patterns of birds were routinely deprioritized due to hurricane or severe weather modeling.

Single model would take 3-4 weeks to complete, an entire species migration would take months

Researchers had many species to model to make a decision on the windmills.

Azure: More Questions Answered Faster

NOAA had used Azure only for customer facing websites.

Enter Azure for Big Compute workloads:

- No Queues / Prioritization
- New Cost Model (scales linearly on demand)
- Rapid deployments using Azure Resource Templates
- Used DS14v2 with Premium Storage and OSS HPC Scheduler (Altair PBS Pro)

POC – 2.5 hours for a single species

Scaled up over the next week to many simulations, and now running in production for other research workloads.

"There was a job that took nine months to run in house that took three days on Azure," Robert Aspinall, IT team lead

Microsoft Azure for Research

Microsoft Research

Apply for a Microsoft Azure for Research Award

We're offering grants of Microsoft Azure to university and non-profit research labs.

Microsoft Research is soliciting proposals for the use of Microsoft Azure in research. We welcome research proposals from any branch of scholarly activity. We are interested in individual investigator projects as well as projects that will support access to services and data of value to a collaboration or community. In addition, we will periodically announce additional special-opportunity RFPs on specific cloud research topics. Winning proposals will be awarded allocations of Microsoft Azure storage and compute resources for a period of one year. The size of the allocations will be suitable for a substantial research project.

Qualified applicants must be affiliated with a university or non-profit research organization. Your proposal should be in English and less than three pages in length. It should include resource requirement estimates (number of core, storage requirements, and so forth) for your project.

To read about current RFPs and submission deadlines, please visit Award Program. For more information, questions, or to make changes to your submitted proposal, please send email to azurerfp@microsoft.com.

The information you provide will be used to verify eligibility and to improve and personalize your relationship with Microsoft Research. Microsoft may contact you as part of the verification process. Microsoft respects your privacy. To learn more, please read our Privacy Statement.

Learn more

- Azure for Research Award program
- FAQ about Azure for Research Awards
- Proposal Eligibility

Learn more

What it is

Azure cloud computing helps you accelerate your research by providing what you need, when you need it. The Microsoft Azure for Research program awards cloud computing time, training, and resources to help you achieve more.

Who it's for

Faculty, researchers, and graduate students are qualified to submit proposals for Azure awards for research projects. Masters and undergraduate students require a faculty project supervisor to submit their proposal.

Apply

Apply now for cloud computing resources for your research project. Proposals are evaluated every two months.

We are committed to the Research community: Azure Data-Egress Fee Waiver

We want the cloud to work the way the academic research community works. That's why we're waiving the Internet egress fee for qualified customers.

- Ensures predictability and stability in cloud costs
- Removes a significant barrier to cloud adoption, allowing researchers to move data freely
- Paves the way for researchers to accelerate the pace of the important work they're doing

Internet2 GÉANT Jisc SURFnet HEAnet Ireland's National Education & Research Network

Accelerate genomic processing on Microsoft Azure

Microsoft has developed a much faster way to do a key part of genomic sequencing with an Azure-powered version of Broad Institute's Genome Analysis Toolkit...

Learn more

Microsoft Genomics

The demand for cloud computing solutions is clear. Today's medical centers, integrated delivery networks and labs are looking for agility, easier management, and access to more capacity to enable the increased demand for next-generation sequencing (NGS). Health and Life Science organizations can benefit from Microsoft's industry-leading approach to security, privacy, and local compliance in the cloud while minimizing cost and complexity. Find a partner→

https://notebooks.azure.

Secure | https://notebooks.azure.com

Microsoft Azure Notebooks Preview

Sign In

Libraries FAQ/Support What's New

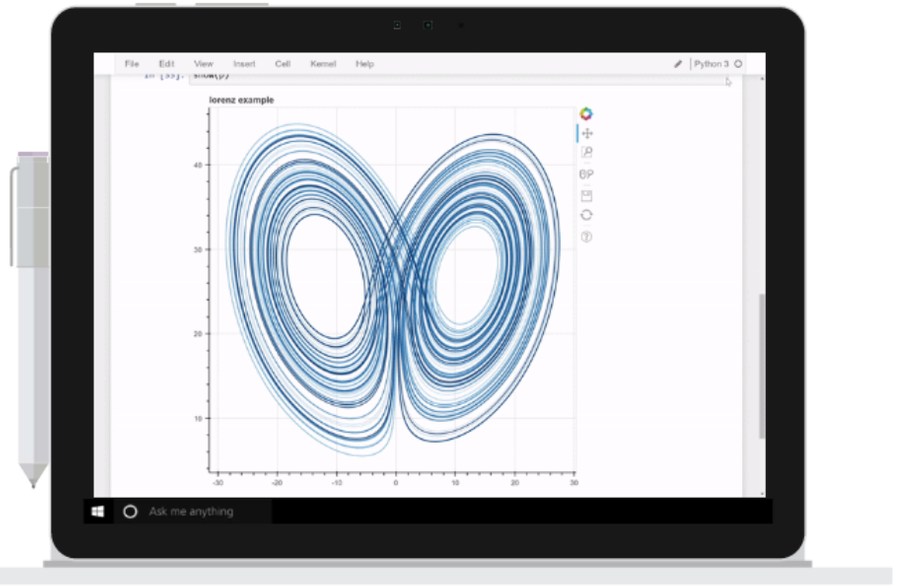

Sharing Your Ideas Made Easy

With Azure Notebooks, unleash your ideas in the cloud with the Jupyter Notebook

Learn more

For Data Science and Analysis

For Teaching and Learning

For Research and Development

Research gravitational waves

Investigate data provided by the Laser Interferometer Gravitational-Wave Observatory (LIGO) and make groundbreaking discoveries.

Thank you!

For more information:
http://microsoft.com/hpc